

AD-A068 428

UTAH UNIV SALT LAKE CITY DEPT OF COMPUTER SCIENCE
APPLICATION OF ADAPTIVE NOISE CANCELLATION TO NOISE REDUCTION I--ETC(U)
MAR 79 D C PULSIPHER
UTEC-CSC-79-022


F/G 20/1

N00173-77-C-0041

NL

UNCLASSIFIED

OF /
ADA
068428



Application of Adaptive Noise Cancellation to Noise Reduction in Audio Signals


DENNIS CARL PULSIPHER

UNIVERSITY OF UTAH

③ LEVEL

AD A068428

DDC FILE COPY



DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

MARCH 1979

UTEC-CS-79-022

COMPUTER SCIENCE, UNIVERSITY OF UTAH
SALT LAKE CITY, UTAH 84112

79 05 08 024

The views and conclusions contained in this document are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency of the U. S. Government.

This document has been approved for public release and sale; its distribution is unlimited.

6

APPLICATION OF ADAPTIVE NOISE CANCELLATION
TO NOISE REDUCTION IN AUDIO SIGNALS

by

10

Dennis Carl Pulsipher

DDC
RECEIVED
MAY 9 1979
B

12 72e

11

March 1979

14

UTEC-CSc-79-022

This research was supported by the Advanced Research Projects Agency of
the Department of Defense under Contract No. N00173-77-C-0041

15

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

✓
404 949

JLB

ABSTRACT

The LMS Adaptive Noise Cancellation algorithm has been applied to the removal of high-level white noise from audio signals. Simulations and actual acoustically recorded signals have been processed successfully, with excellent agreement between the results obtained from simulations and the results obtained with acoustically produced data. A study of the filter length required in order to achieve a desired noise reduction level in a hard-walled room is presented. The performance of the algorithm in this application is described and required modifications are suggested.

A multichannel processing scheme is presented which allows the adaptive filter to converge at independent rates in different frequency bands. This is shown to be of particular use when the interfering noise is not white. Careful implementation of the scheme allows the problem to be broken into several smaller ones which can be handled by independent processors, thus allowing longer filter lengths to be processed in real time.

ACCESSION for	
NTIS	White Section <input checked="" type="checkbox"/>
DDC	Buff Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION _____	
BY _____	
DISTRIBUTION/AVAILABILITY CODES	
Dist.	AVAIL. and/or SPECIAL
A	

TABLE OF CONTENTS

ABSTRACT	iv
LIST OF FIGURES	vii
CHAPTER 1 INTRODUCTION	1
Research Objectives	1
Content Review	2
Historical Context	3
Literature Review	7
CHAPTER 2 SYSTEM DESCRIPTION	12
Optimization Criterion	12
Block Solution	13
Adaptive Solution	17
Data Generation	20
CHAPTER 3 PRACTICAL CONSIDERATIONS	23
Real Time Performance	23
Multichannel Processing	25
Other Suggested Modifications	29
Further Problem Aspects	31
CHAPTER 4 EXPERIMENT DESCRIPTION	35
Basic Experiments	35
Noise Reduction Measurements	36
Synthetic Experiments	38

Room Simulations	42
Multichannel Experiments	43
Acoustically Recorded Experiments	47
CHAPTER 5 CONCLUSIONS	50
Observations on Results	50
Summary	54
Future Work	55
APPENDIX A CALCULATION OF THE CONDITIONAL MEAN	56
APPENDIX B CHARACTERISTICS OF THE LMS ALGORITHM AND NOISE CANCELLATION	59
REFERENCES	62

LIST OF FIGURES

1.1	LINEAR FILTERING	5
1.2	MODEL FITTING	6
1.3	NOISE CANCELLING	8
2.1	BASIC DATA GENERATION	21
2.2	SIMPLE DATA GENERATION MODEL	22
2.3	EQUIVALENT DATA GENERATION MODEL	22
3.1	BLOCK DIAGRAM OF MULTICHANNEL PROCESS	26
4.1	TWO SYNTHETIC CHANNELS	37
4.2	A MULTILEVEL SYNTHETIC CHANNEL (G_2)	38
4.3	AN EXAMPLE OF SIGNALS PROCESSED	40
4.4	SPECTRA OF THE SIGNALS IN FIGURE 4.3	41
4.5	TYPICAL ROOM CHANNELS	42
4.6	A CHANNEL AND THE ESTIMATES OF ITS INVERSE	45
4.7	SPECTRA OF CHANNELS IN FIGURE 4.6	46
4.8	RESULTS OF ACOUSTICALLY RECORDED EXPERIMENTS	49

CHAPTER 1

INTRODUCTION

Research Objectives

In this research, Adaptive Noise Cancellation (ANC) has been applied to the enhancement of speech signals in environments where the noise corrupting a signal contains as much energy as the speech itself, or more. The bulk of the research performed has had as its objective the evaluation of the algorithm's performance for this application. In order to evaluate this performance, the problem has been approached in four stages. The first stage was the creation of completely controlled synthetic situations and application of the technique. The second was application of the technique to actual situations. The third stage was the evaluation of problems noted, and searching for feasible explanations, with verification at each stage. And finally, suggestion and testing of solutions to the encountered difficulties was performed. For more information about the actual experiments themselves, the reader is referred to chapter 4.

In following this approach, the following research objectives have been achieved:

- 1) Adaptive Noise Cancellation has been successfully applied to the removal of high level white noise from noisy speech signals for the first time.
- 2) A technique capable of enhancing speech in extremely noisy environments has been made available to speech processors.
- 3) A tool capable of adaptively removing non-stationary noise has been applied to speech problems.
- 4) A "multichannel" modification of ANC which improves its performance by improving its convergence characteristics has been developed and implemented successfully.
- 5) A quantification of results which can be expected from ANC under varying conditions has been made.

Content Review

In this research, noise cancellation techniques have been applied to audio signals. The goal has been to improve the intelligibility of speech produced in an acoustically hostile environment. It has been assumed that this need has arisen because of a desire to encode the speech for low bandwidth digital transmission.

Chapter 1 continues with a description of the problem being addressed and a review of previous work. Chapter 2 then describes the noise cancelling system, along with its mathematical foundations, and the noise generation model to which it is particularly applicable. In chapter 3,

many practical aspects of the problem are discussed and "multichannel processing" is described. This technique is proposed as a solution to several of the problems encountered in the implementation of noise cancellation for actual use. A description of the experiments performed and their results follows in chapter 4. The interpretation of the results is reserved for chapter 5. Related developments and descriptions of some properties of the noise cancellation algorithm are found in the appendices.

Historical Context

In recent years military cryptographers and others interested in secure communication channels have placed increased importance on digital communication links. In order to meet low-bandwidth requirements, several digital encoding schemes have been developed. These schemes, such as Linear Predictive Coding (LPC), take advantage of the characteristics of speech, making possible low-bandwidth digital transmission of speech signals. Unfortunately, these schemes are typically sensitive to noise in the environment. While many noise-reduction schemes have been tried, none have proved successful in extremely noisy environments or in environments with nonstationary noise sources. There is a need for the development of techniques which can provide noise reduction in these situations. This research concerns the application of a

noise cancellation technique which is capable of working in extremely noise environments and powerful enough to deal with certain kinds of non-stationarity in the environment.

While noise reduction has long been desirable to minimize the problems mentioned, it has taken on increased importance in recent years. The advent of low-cost, high-speed digital hardware has made possible the use of complex digital encoding schemes for transmitting speech digitally in low-bandwidth channels. While such vocoders perform in quiet environments, they often fail miserably in acoustically hostile environments. The sources of these failures can often be classified into four categories: silence detection, spectral estimation, voiced-unvoiced determination, and pitch period estimation [1]. Unfortunately a scheme which improves performance in one of these areas, does not necessarily aid in the others.

Noise reduction techniques have generally fallen into one of three categories: linear filtering, model fitting, and noise cancellation. The traditional approach has been some sort of linear filtering, while model fitting has been used less often, and noise cancellation has been used even less frequently.

Many methods have been proposed for determining linear filters for removal of different types of noise. All of these attempt to improve the signal-to-noise ratio

by attenuating frequencies where the signal-to-noise ratio is poor and emphasizing those frequencies where the signal-to-noise ratio is better. The pioneering work of Wiener [2] resulted in perhaps the most often used example of these methods. Kalman filtering [3] is also included in this group, and demonstrates the capability of these techniques to handle nonstationary processes. A drawback to using linear filtering techniques for noise reduction is that the desired signal is also passed through the filter and is thereby modified (see fig. 1.1). Distressingly, signals whose quality has been improved by such methods do not always retain an intelligibility improvement when subsequently digitally encoded for low bit rate transmission [4].




Figure 1.1 LINEAR FILTERING

Model fitting methods use a priori information about the class of signals to be transmitted. Parameters describing the signal are determined and sent. The signal estimate is then synthesized from the estimated parameters. Transmission of a single digital bit is such a scheme

which has found wide use because of its inherent noise reduction capability. Linear Predictive Coding is another such technique which has found wide use in the speech processing community. In the field of speech encoding it is perhaps the best known of these methods, though it is not particularly effective for noise reduction. The Homomorphic vocoder developed by Neil Miller at the University of Utah [5] is another example of model fitting techniques.

Noise reduction by model fitting suffers from two major problems. Developing a model which is not sensitive to noise, is a formidable, if not impossible task. And when an acceptable model is derived, it is still necessary to estimate its parameters accurately if noise reduction is to be achieved. In the presence of noise, such parameter estimation is not a trivial matter. In fact, it is in this area that many of the often used coding techniques such as LPC appear to fail in noisy environments [6] (see fig. 1.2).




Figure 1.2 MODEL FITTING

Noise-cancellation methods are significantly different in philosophy from other noise-reduction techniques. They attempt to estimate the actual disturbing noise in the time domain and perform an algebraic subtraction of the noise estimate to produce a signal estimate. These procedures have been used for antenna side-lobe cancellation [7]-[8], digital channel equalization [9], telephone channel echo cancellation [10]-[13], noise reduction in electro-cardiography [14], and spectral line enhancement [15]. An especially desirable attribute of these methods is that if the noise has been generated according to the modelled form, the signal can be recovered unaltered. In practice, the signal estimates obtained from such techniques are compatible with LPC methods. This is because the signal's phase is not altered by passing through a filter (see fig. 1.3). They do, however, require a noise reference channel and, while being conceptually simple, they are often computationally demanding.

Literature Review

Before describing the work completed during this research, let us review related earlier work. In doing this it is well to remember that much of the renewed emphasis in noise reduction in speech signals is a result of the development of the Linear Predictive Coding scheme by Atal and Hanauer [16]. McAulay [1] and Makhoul [17], among others, pointed out the less than desirable

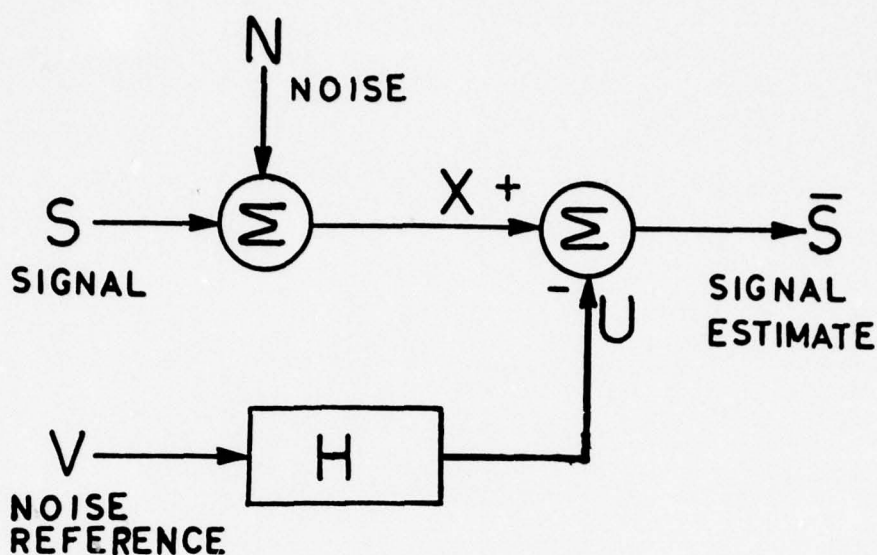


Figure 1.3 NOISE CANCELLING

performance achieved by this scheme in the presence of noise. Yegnanarayana described the causes of some of these failures, as well as criticized possible solutions to the problems [6].

Many techniques were then applied to the problem of improving speech intelligibility in noisy environments. These techniques have fallen into two broad categories: preprocessing techniques, which attempt to remove the noise before spectral coding takes place, and algorithm modifications, which attempt to integrate noise reduction with the encoding process itself.

Most of the preprocessing work that has been done has been linear filtering for noise removal. Frequently this

has been along the lines suggested by Wiener [2]. A review of linear filtering methods and developments of recent years has been published by Kailath [18]. Included in that paper is an extensive bibliography to which the reader is referred for more linear filtering references. A small amount of work has been done in defining models to be used as preprocessors; Boll's SABER algorithm is such a development [19].

A majority of the algorithm modifications which have been proposed have been modelling changes to incorporate new capabilities or refine old ones. Boll [20], with Predictive Noise Cancellation, proposed a method of modelling the noise and, hence, minimizing its effects. Atashroo [21] proposed a method of pole-zero modelling to get a better spectral fit. And Christiansen, in a word recognition context, proposed modifying the autocorrelations of the noisy signal under the assumption that the signal and noise were uncorrelated [22]. He then concluded that such a scheme was not promising since the assumption, while true statistically and for long time averages, was almost never true over the short time windows used in Linear Predictive Coding analysis.

Only recently has much work been put into algorithm modifications to increase the accuracy of parameter estimates. This work has centered around stochastic estimation techniques. The reader is referred to an early

paper by Bode and Shannon [23] and a more recent text by Box and Jenkins [24] for more information about modelling and parameter estimation.

In this research, Adaptive Noise Cancellation has been applied as a pre-filtering technique of noise reduction. This technique uses an algorithm related to one proposed by Robbins and Monroe [25] and analyzed by Sakrison [26] to remove the noise from the noisy signal algebraically. Its basic form was developed by Widrow and Hoff [27], and described by Widrow [28], Senne [29], Kaunitz [14], and Widrow, et al. [30].

Senne [29] analyzed the behavior of the algorithm under assumptions of Gaussian inputs and independent reference noise measurements. Daniell [31] extended the convergence proofs to include certain types of correlated noise measurements, and Kim and Davisson analyzed the effects of "M-dependence" and requirements to guarantee convergence under such conditions [32]. Kaunitz [14] extended the conditions under which the algorithm converges by implementing a random reference noise sampler for the purpose of updating the filter. Widrow, et al. [33] also analyzed the algorithm's performance under certain nonstationary conditions.

A simpler algorithm was proposed by Moschner [34] and researchers at Bell Laboratories [12] which was more easily implemented but had less favorable convergence

properties. McSherry [35] and Nagumo and Noda [36] suggested the use of a slightly different gradient search approach which allowed the algorithm to maintain a constant adaptation "time-constant" even though its energy fluctuated widely. Gitlin, Mazo, and Taylor [37] discussed the design of gradient algorithms for digital applications, and Frost [38] described an algorithm which was adaptive, but subject to equality constraints.

Noise cancellation has been applied to many problems in the past. Riegler and Compton [7] and Widrow, et al. [38] applied the technique to antenna interference rejection. Lucky [9] suggested its use for digital channel equalization, while Sondhi [10], Mueller [11], Rosenberger and Thomas [13], and Weinstein [12] have described its use for echo cancellation in the telephone network. Glover [15] described its use for the extraction of narrow band signals or noise. And Kaunitz [14] performed experiments using it for noise reduction in electrocardiography. He also performed some experiments wherein highly stylized noise was removed from a noisy speech signal which was meant to simulate an aircraft cockpit. Widrow, et al. [30] reported many of these results and uses in a review of noise cancelling developments late in 1975.

CHAPTER 2

SYSTEM DESCRIPTION

Optimization Criterion

Noise cancellation is achieved by algebraically subtracting a noise estimate from the current noisy signal. Since this could easily result in an increase in noise power at the output of the system, rather than the desired decrease, we ought to examine the mechanism by which this is avoided.

Let us assume that we are given $x(t)$, the sum of two mutually uncorrelated signals, $s(t)$ and $n(t)$, and a third signal $v(t)$, which is mutually uncorrelated with $s(t)$. We can then form a signal estimate

$$(2.1) \quad \bar{s}(t) = x(t) - u(t) = s(t) + [n(t) - u(t)],$$

where $u(t)$ is a noise estimate which we will constrain to be a linearly filtered version of $v(t)$.

Then

$$(2.2) \quad \bar{s}(t)^2 = s(t)^2 + 2s(t) [n(t) - u(t)] + [n(t) - u(t)]^2$$

and

$$(2.3) \quad E\{\bar{s}(t)^2\} = E\{s(t)^2\} + 2E\{s(t) [n(t) - u(t)]\} + E\{[n(t) - u(t)]^2\}.$$

Since $s(t)$ is uncorrelated with $v(t)$ (and hence $u(t)$),

$$(2.4) \quad E\{\bar{s}(t)^2\} = E\{s(t)^2\} + E\{[n(t) - u(t)]^2\} .$$

The mean power of the signal estimate is the sum of the mean power of the signal and the mean power of the noise estimation error $[n(t) - u(t)]$.

Since the mean power of the signal is fixed, minimizing the mean output power minimizes the power in the noise estimation error, which is equal to the power in the signal estimation error. Therefore, minimizing the mean output power causes the signal estimate $s(t)$ to be a least mean squares fit to the signal $s(t)$. The minimization, of course, must be carried out by choosing an $h(t)$ (the impulse response of the filter through which $v(t)$ is passed to generate $u(t)$) which minimizes the power in $s(t)$. We, then, are looking for $h(t)$ which satisfies

$$\underset{h(t)}{\text{Min}}[E\{\bar{s}(t)^2\}] .$$

In practical situations additional uncorrelated noises may be present. For development details of such a situation refer to Appendix A, where the conditional mean of such a process is calculated.

Block Solution

Since it is intended to implement noise cancellation using a digital filter, it seems that a solution for $h(t)$

in vector notation is now appropriate. Let v_n , x_n , s_n , etc. be the value of the corresponding signal at time nT , where T is the sampling interval. (We implicitly assume a band-limited signal at this point.)

Define the vectors

$$(2.5) \quad V_n = \begin{bmatrix} v_n \\ v_{n-1} \\ . \\ . \\ v_{n-L+2} \\ v_{n-L+1} \end{bmatrix} \quad H_n = \begin{bmatrix} h_{1,n} \\ h_{2,n} \\ . \\ . \\ h_{L-1,n} \\ h_{L,n} \end{bmatrix} ,$$

where L is the length of the filter to be estimated and H_n is the filter. We then have

$$(2.6) \quad u_n = V_n^T H_n = H_n^T V_n .$$

The sample signal estimate is calculated by subtracting u_n from x_n .

$$(2.7) \quad \bar{s}_n = x_n - u_n = x_n - V_n^T H_n = x_n - H_n^T V_n$$

Squaring yields

$$(2.8) \quad \bar{s}_n^2 = (x_n - u_n)^2 \\ = x_n^2 - 2x_n V_n^T H_n + H_n^T V_n V_n^T H_n .$$

And taking the expected value gives

$$(2.9) \quad E\{\bar{s}_n^2\} = E\{x_n^2\} - 2E\{x_n V_n^T H_n\} \\ + E\{H_n^T V_n V_n^T H_n\} .$$

Assuming a stationary channel H gives

$$(2.10) \quad E\{\bar{s}_n^2\} = E\{x_n^2\} - E\{x_n V_n^T\} H + H^T E\{V_n V_n^T\} H .$$

Defining

$$(2.11) \quad P = E\{x_n V_n\}$$

and

$$(2.12) \quad R = E\{V_n V_n^T\}$$

we have, at this point assumed that P and R are not functions of time, yielding

$$(2.13) \quad E\{\bar{s}_n^2\} = E\{x_n^2\} - 2P^T H + H^T R H$$

which is a quadratic function of H , hence has a unique minimum H^* . Because of our assumptions of stationarity this filter will also be stationary. By differentiating with respect to the elements of H we get

$$(2.14) \quad \nabla = - 2P + 2R H .$$

Setting $\nabla = 0$ to find the optimal H , we get

$$(2.15) \quad H^* = R^{-1} P .$$

In reviewing the implications of this result, we should remember that it is quite possible for H to be a stationary channel, while $x(t)$ and $v(t)$ are not stationary processes. It is the product of the inverse of the auto-covariance matrix and the cross-correlation vector which must be stationary. If both change with time in such a manner that their product remains constant, the optimal channel will be stationary. It is this condition, in practice, which is required to be nearly true.

In order to calculate this optimal filter, an estimate of the auto-correlation matrix R and the cross-correlation vector P must be made. The necessity of inverting R precludes allowing the time response of the filter to have extremely large numbers of active (non-zero) points. We must be satisfied with a finite length filter, albeit an optimal finite filter. We must, through knowledge of the process, or by arbitrary decree, specify the length of the filter, and the amount of delay to be incorporated to allow the creation of an effectively non-causal filter. The specification of these parameters can be thought of as specifying the "active interval" or "domain" of the filter.

In performing the experiments described as "block analyses," the active intervals of the filters were chosen to correspond with the active intervals chosen for the "adaptive analyses." The filters themselves were

calculated using a standard Levinson's recursion algorithm [40] operating on auto- and cross-correlation estimates which were calculated by averaging sample correlations of shorter "blocks." The sample correlations were calculated by taking straightforward inner products of appropriate data vectors. All such inner products were of the same length for any given experiment, and no zeroes were appended to any data vector. That is, the sample correlations were not calculated as "windowed" correlations, but as true sample correlations. Succeeding "blocks" were disjoint, and the union of all used blocks was the entire signal set. Thus, the final optimal filter estimate is a globally optimal finite filter for the active interval specified.

Adaptive Solution

In practice, the channel to be estimated cannot always be guaranteed to be stationary. For this reason it is felt that a non-terminating adaptive estimator is most applicable. That is, we want to use an estimator which continues to adapt even after it has achieved a good channel estimate.

Since we want to calculate H adaptively we might try a standard steepest descent algorithm

$$(2.16) \quad H_{n+1} = H_n - \mu \nabla_n$$

where the parameter μ controls convergence and stability.

Unfortunately, we do not have access to V_n , so must be satisfied with a gradient estimate \hat{V}_n . Widrow [28] has suggested the use of

$$(2.17) \quad \hat{V}_n = -2\bar{s}_n V_n$$

which yields the algorithm

$$(2.18) \quad H_{n+1} = H_n + 2\mu\bar{s}_n V_n .$$

Others have suggested many similar, but slightly different algorithms to which we have referred in chapter 1.

By defining the expected value of H as M it is simple to see that

$$(2.19) \quad M_n = [I - 2\mu R]^n H_0 + R^{-1}P - [I - 2\mu R]^n R^{-1}P .$$

By diagonalizing R , it is a short step to show that

$$(2.20) \quad \lim_{n \rightarrow \infty} \{M_n\} = R^{-1}P \quad \text{for } 0 < \mu < \frac{1}{\lambda_{\max}}$$

where λ_{\max} is the largest eigenvalue of the matrix R . The variance of the estimate can also be forced below any arbitrary positive limit as n gets large for V_k uncorrelated with V_j for $k \neq j$ [29]. Under further assumptions convergence has also been shown for special cases of correlated V_j . Asymptotic behavior, residual error, and nonstationary behavior in special cases have also been investigated elsewhere.

Since the algorithm converges to $R^{-1}P$, it is of some interest to look briefly at its characteristics under conditions similar to those used in calculating the conditional expectation of S in Appendix A. That is, let the reference noise V be the sum of two mutually uncorrelated random variables. One of these (N) can be correlated with the noise in the noisy signal, while the second ($M2$) is to be uncorrelated with the noisy signal. Then

$$(2.21) \quad V = N + M2$$

and

$$\begin{aligned} (2.22) \quad R &= E\{V_n V_n^T\} = E\{(N_n + M2_n)(N_n + M2_n)^T\} \\ &= E\{N_n N_n^T + N_n M2_n^T + M2_n N_n^T + M2_n M2_n^T\} \\ &= R_{NN} + R_{M2M2} \end{aligned}$$

so

$$\begin{aligned} (2.23) \quad R^{-1}P &= [R_{NN} + R_{M2M2}]^{-1}P \\ &= R_{NN}^{-1} \{R_{NN}[R_{NN} + R_{M2M2}]^{-1}\}P \end{aligned}$$

The braced terms show the Wiener Filtering performed to obtain N from V , while the remaining terms perform the transformation mentioned in Appendix A. Since it is now being performed without a priori knowledge of the statistical behavior of $x(t)$ and $v(t)$, only the subsequent

subtraction is performed, leaving the final filtering described in the appendix undone. Of course, if statistics are known, the final filtering operation should be performed.

Since the optimal filter is a function of the inverse of R , $R^{-1}P$ (eq. 2.15), it seems appropriate to consider the conditioning of R . If R is singular, additional conditions must be imposed to obtain an optimal solution [39]. This does not mean, in general, that there is no solution, simply that it is not unique. This condition is frequently encountered in practice when the interfering noise is periodic, or nearly periodic. While channel estimation is not completely possible in such cases, it is only necessary to estimate the channel accurately in those frequency bands where significant interfering energy is present. Even though the channel estimate may be considered poor in such a situation, the noise reduction achievable may be significant.

Data Generation

The manner of data generation is irrelevant to the preceding developments. But, by suggesting a data generation scheme, greater appreciation for the noise cancellation procedure can be fostered and difficulties more easily understood. It is clear that if the data can be accurately modelled as shown in figure 2.1, and if the channel can be accurately modelled as a finite length all-

zero filter, perfect noise cancellation can be achieved if the estimated linear filter, H , is set equal to G .




Figure 2.1 BASIC DATA GENERATION MODEL

Figure 2.2 shows a model which resembles actual acoustic signal generation more closely. Its equivalence to the previous model (as shown in figure 2.3) gives us hope that a high degree of noise cancellation is still possible.

Additional noise sources, such as extra uncorrelated noises at the noisy signal and noise reference pick-ups, degrade performance, but are easily analyzed, as was done above. If these extra sources happen to be mutually

correlated, the channel estimator is faced with the task of estimating a combination of two separate channels. Since no such equivalent channel exists, the amount of noise reduction which can be achieved is highly dependent upon the signals involved.




Figure 2.2 SIMPLE DATA GENERATION MODEL




Figure 2.3 EQUIVALENT DATA GENERATION MODEL

CHAPTER 3

PRACTICAL CONSIDERATIONS

Real Time Performance

Since an objective of this research has been to apply a relatively untried procedure to the problem of noise reduction in acoustically hostile environments, the rate at which the procedure can operate is of considerable importance. If the adaptive filter is updated at the sampling rate, we must perform two multiplies and additions for each point in the active interval of the filter during each sampling interval. One of these multiplies and additions is performed to produce the output of the adaptive filter, and the other is performed in the process of updating the filter itself. Of course there are a few additional operations which can be viewed as overhead, since they do not depend on the filter length. If the number of points in the filter's active interval is large, this task is formidable. In order to accomplish it, some type of parallel processing may need to be performed.

The most straightforward application of parallel processing could be called time-domain parallelism. In applying this procedure the troublesome calculations are broken into sets of smaller calculations which can

be performed at the necessary rate by several processors working simultaneously.

The troublesome operations are

$$(3.1) \quad u_n = H_n^T V_n$$

and

$$H_{n+1} = H_n + 2\mu \bar{s}_n V_n .$$

Let us partition H_n and V_n , so that

$$H_n = [H_{1n}^T H_{2n}^T]^T$$

and

$$V_n = [V_{1n}^T V_{2n}^T]^T .$$

We see that equation 3.1 may be rewritten as

$$u_n = u_1 + u_2 .$$

We also see that

$$H_{1n+1} = H_{1n} + 2\mu \bar{s}_n V_{1n} ,$$

while

$$H_{2n+1} = H_{2n} + 2\mu \bar{s}_n V_{2n} ,$$

where

$$u_1 = H_{1n}^T V_{1n}$$

and

$$u_2 = H_{2n}^T V_{2n} .$$

The problem has been broken into two smaller problems which can be handled by separate arithmetic units.

Further division could, of course, be performed until the

problem has been reduced to a size which can be performed by available units working in parallel.

A second application of parallel processing techniques could be viewed as frequency division parallelism of multi-channel processing, though all of the processing is performed in the time domain (see fig. 3.1). In this scheme, the reference noise is passed thorough a bank of band-pass filters. It is preferable if these filters are of equal bandwidths occupying separate frequency bands, and having the same linear phase. The bands, when considered as a whole, should span the part of the spectrum occupied by the original reference noise. These constraints simplify implementation and visualization of the process, but are not constraints arising from the mathematical description of the procedure. Besides allowing parallel processing, this process performs significantly better than the standard ANC technique in some situations where the adaptive filter fails to converger properly when implemented in the standard manner. A further description is found in the following subsection.

Multichannel Processing

Multichannel processing is a procedure which allows parallel processing of noisy signals. In addition, it improves the convergence properties of the noise canceling algorithm. If the reference noise itself is not white, a severe problem may be encountered while employing

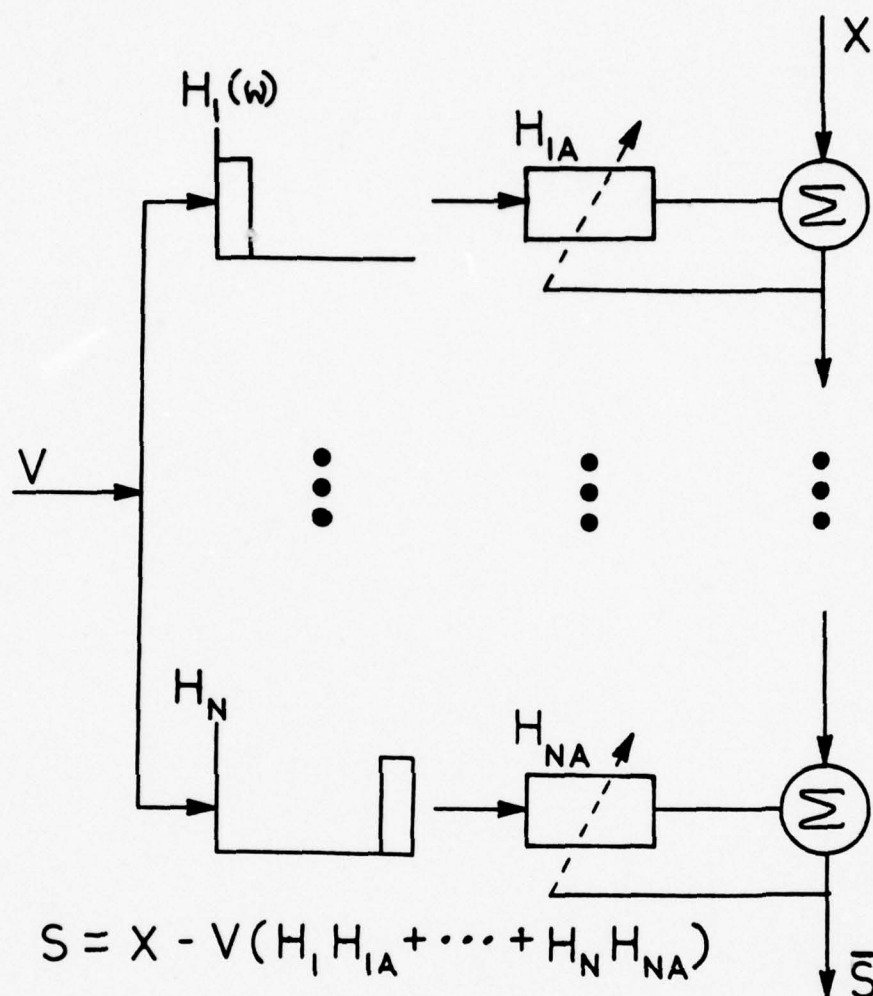


Figure 3.1 BLOCK DIAGRAM OF MULTICHANNEL PROCESS

the standard ANC algorithm. Convergence is not necessarily guaranteed in such cases, and if it does take place, the rate at which it occurs is not easily estimated. If the noise has varying spectral densities in different frequency bands, it will converge at different rates in the different bands, if it converges at all.

This is because the stability of convergence and convergence rate are related to the energy in the noise reference signal. For a more detailed description of these properties the reader is referred to appendix B. If an attempt is made to increase the adaptation rate in order to achieve a desired convergence rate in a band with low energy, the estimation process may go unstable; the signal estimation error will be increased.

This seemingly difficult problem can be dealt with in two ways. First, it can be ignored. If the noisy signal does not contain significant energy in the frequency bands where the reference noise has low energy, it may not be necessary to have an accurate channel estimate in those bands; after all, it is noise reduction, not channel estimation, which we desire. And second, the signal may be treated as "piecewise white," that is, it may be divided into frequency bands which contain nearly constant spectral densities. These bands may be used as independent noise reference inputs and multi-channel processing employed. Because the different frequency bands can then be estimated by independent noise cancellers, the convergence rates of each band may be specified independently, thus allowing all bands to converge at the same rate (or different rates), even though the reference energy in each band is different.

A description of the basic process, shown in figure 3.1, follows. The output of one of the band-pass filters is applied to a noise cancelling processor as the noise reference input, while the noisy signal is applied as the other input; the resulting signal estimate is then applied as the noisy signal to a second, independent noise cancelling processor, which uses as its noise reference input, the output of a second of the band-pass filters. The resulting signal estimate can be subsequently treated as a noisy signal, and the procedure continued by putting in tandem further independent processors until all of the available reference inputs have been used.

If the process is implemented directly as described, convergence benefits may be derived, but the number of calculations required to be performed by each processor will be increased. Since the active interval of the filter remains constant, the number of calculations required is proportional to the number of processors performing them. Thus, each processor must perform as many calculations as if it were doing the job alone. If the bandwidth of each band-pass filter is constrained to be less than a factor of N smaller than the bandwidth allowed by the sampling rate, however, the number of operations required of each processor can be reduced by a factor of N . This can be achieved by effectively down-sampling the noise reference inputs by a factor of N ,

since their bandwidths allow it. This down-sampling can be performed without the need for a complicated interpolation scheme at the filter output if it is done by simply allowing only one filter point in N to be non-zero. Only those points, then, are used in performing the necessary calculations. While this causes frequency-domain aliasing of the filter itself (it appears to be under-sampled with respect to the original sampling rate), the incoming reference data has already been filtered to prevent it from containing spectral components in any but one of the reduced bandwidth spectral copies. Each of the independent noise cancellers can thus be required to perform a smaller number of calculations than that required if the job were being done by a single processor.

Other Suggested Modifications

In examining the performance characteristics of the basic noise cancellation algorithms described in chapter 2, it has become apparent that certain deficiencies exist. By altering the algorithms slightly some of these deficiencies can be minimized. Among the properties requiring improvement are the ability to deal with nonstationary input signals, convergence, and convergence rate stability.

If the block analysis approach is implemented, rather than an adaptive approach, a large delay is introduced into the system. Furthermore, if the signals are not stationary, block analysis is not likely to produce any

sort of optimal filter for any particular segment of the signal. It attempts to generate a filter which results in the minimum total output power, but may do a very poor job on some segments of the signal. In short, it was not developed for nonstationary signals. By block processing short segments, however, it could be adapted for use with some kinds of nonstationary signals.

If the adaptive procedure is implemented, the primary difficulties center on the convergence properties of the algorithm. Previous developments of convergence rate and signal estimation error have relied on the assumption that the eigenvalues of the noise reference auto-covariance matrix are constant and equal (see appendix B). This is related directly to requirements of stationarity of the reference noise. It relates, in particular, to the assumptions that the reference noise has a constant energy, and is white. In practice, these assumptions could be quite unfounded. The consequences of the violation of these assumptions are quite direct. The convergence rate of the adaptive filter becomes time-varying, causing possible instabilities, and the signal estimation accuracy also becomes time-varying. Since this problem is caused by the dependence of convergence on the eigenvalues of R (and hence the energy of the noise reference signal) it can be alleviated by implementing some sort of automatic gain control. This

may be introduced in both input signal paths and the output signal path, or in the feedback path alone. A simple solution has been described by McSherry [35] and Nagumo and Noda [36]. It results in equation (2.18) being modified to be

$$H_{n+1} = H_n + \frac{2\mu}{\text{TR}[V_n V_n^T]} \bar{s}_n V_n \quad .$$

Further Problem Aspects

It should be recognized that there are many other facets of noise cancellation which must be considered in order to attain the degree of performance desired. While some of these may seem trivial, and others beyond our control, we should at least be aware of problems that exist.

A knowledge of the type of noise to be removed is essential. While details may be unimportant, the general characteristics of the noise must be known before a noise cancellation strategy is specified. For example, if the noise is periodic, or nearly periodic, an entirely different noise cancellation strategy will be employed than that described for the removal of broad-band noise. The adaptation rates and filter lengths specified will undoubtedly be different in the two cases. Periodic noise can be eliminated with a much shorter filter than can white noise, the adaptation rate can be made correspondingly

faster with no loss in accuracy, and the noise reduction can often be accomplished without an independent reference noise measurement. For more details concerning the removal of narrow-band noise the reader is referred to Glover [15].

Since the noise cancelling filter is to be implemented as a transversal digital filter, several additional constraints are implied. First, by the mere act of specifying digital processing we have implied that many conditions will be met. The signals to be digitized will be band-limited to a frequency corresponding to half the sampling rate; a certain amount of quantization error will be allowed; the sampling will be done accurately at the ends of uniform time intervals. Second, by using a transversal filter, we have constrained the filter to be causal and finite. It is to be a causal and truncated Wiener Filter. Of course, introduction of appropriate delays can result in a filter which is effectively non-causal, but still finite. The determination of such delays is not a well defined procedure, but must be done experimentally for any given application, though they can often be arbitrarily specified because of an understanding of the physics of a particular problem.

While a truly optimal filter may exist, which is independent of noise characteristics, when the filter is constrained to be of finite length, no such filter may

exist. The optimal filter, under such a constraint, may be very dependent on the type of noise present.

Other effects of using a finite length filter depend largely on the environment. If the transfer functions of the environment which transform the noise produced by the source into the noise interfering with the signal at the signal sensor, and the reference noise picked up by the reference sensor, can be modelled as linear filters, the problem is reduced to estimation of $G_2^{-1}G_1$ (see fig. 2.3).

Though they are actually characterized by many poles and zeroes, if G_1 and G_2 (fig. 2.2) can be modelled as all zero filters, as can probably be expected, the difficulty in estimating the optimal filter arises because of the need to effectively invert G_2 . If, as hoped, it was an all zero (Moving Average) process, its inverse will be an all-pole (Auto-Regressive) process. It is highly unlikely that G_2 is a "minimum phase" process, since it is a physically realized process which will undoubtedly introduce phase dispersion. It will have zeroes outside the unit circle in the z -plane. Its inverse will, therefore, have poles outside the unit circle, apparently threatening to cause the inverse to be unstable. This can be avoided, however, by recognizing that such a system corresponds to an impulse response which is stable, but doubly infinite in length. It must be non-causal and

non-truncated. Of course if these singularities are well away from the unit circle, the response will be dominated by rapidly decaying exponential envelopes. As they approach zero we may choose to truncate them and use only those points where the response has had more significant energy. This allows us to approximate the required doubly infinite recursively generated filter, with a finite transversal filter. As the zeroes of G_2 and the actual poles of G_1 approach the unit circle, however, the number of points which we must allow in the active interval of the filter to be estimated grows if we desire to maintain a constant error. The reader is referred to appendix B for more on the subject of filter truncation.

CHAPTER 4

EXPERIMENT DESCRIPTION

Basic Experiments

In order to evaluate Adaptive Noise Cancellation as a technique for noise reduction in audio signals, many experiments were performed. These experiments were performed in four groups. They were designed to identify potential performance, actual performance, problem areas, and evaluate possible solutions.

An experimental data base was created which made possible the comparison of different experimental results. A General Radio Company type 1390-B random noise generator was used as a primary noise source. Its output was low-pass filtered to 3.2 KHz and sampled at a rate of 6.67 KHz for use as a nearly white gaussian noise source which could be used repeatedly. A Hewlett-Packard model 209A oscillator used as a square wave generator was also used to generate a similarly recorded, nearly periodic noise sample. This sample was made highly nonstationary by varying the frequency adjustment of the square-wave generator in a semi-random fashion while digitization was taking place. These samples were then concatenated and used as noise sources for both synthetic and accoustically recorded experiments.

A number of FIR filters were then created in order to allow a group of entirely synthetic experiments to be performed. Five such filters have been used for the majority of the experimental work. A low-pass filter with its cutoff frequency at approximately 1500 Hz and a triple band-pass filter were created to allow an evaluation of the technique's promise (see fig. 4.1). Two room-channel estimates were made from actual measurements of a room's response in order to simulate, digitally, an actual room [41]. And a fifth filter was created by specifying the locations of a set of zeroes in the z-plane. This filter had frequency bands with different, nearly constant gains and was used in studying certain convergence properties of the algorithm (see fig. 4.2).

The noise base samples were passed through these filters and measurements were made to determine the energy contained in these signals before and after filtering. The energies measured later allowed signal-to-noise ratios to be easily specified when digitally creating noisy speech signals by adding scaled versions of these signals to digitally recorded speech signals.

Noise Reduction Measurements

In order to evaluate performance, it was necessary to define a procedure for determination of signal improvement. The noise reduction figure was calculated as the ratio of output power to noisy signal input power averaged over





Figure 4.1 TWO SYNTHETIC CHANNELS




Figure 4.2 A MULTI-LEVEL SYNTHETIC CHANNEL (G_2)

8192 data samples during a period of no speech activity. During this time, the signal could be assumed to be zero, meaning that any input or output signals would be entirely noise. While this is not an acceptable method of noise measurement for many techniques, because of the nature of this technique of noise reduction, it is felt that the method is proper and results representative of its true performance can be thus obtained.

Synthetic Experiments

For the initial experiments digitally recorded speech signals were added to the filtered noise segments and used

as the noisy signals applied to the ANC algorithm, while the corresponding unfiltered noise segments were used as the noise reference input signals. For actual room simulation the noise filtered through one of the room channel estimates was used as the noise reference, while the noise filtered through the other channel was scaled and added to the digitized speech signal for use as the noisy signal. Typical results can be seen in figures 4.3, 4.4 and 4.5.

During these experiments, it was found that perfect cancellation could be achieved during extended periods of no speech activity. Of course, any non-zero signal at the output caused the adaptive feedback mechanism to actively adapt the estimated filter, even though it had achieved a perfect estimate. By referring to equation 2.18, it will be noted that the amount of such degradation is signal dependent. By using a slow adaption rate, however, this degradation can be kept below any desired positive level.

When white gaussian noise was used as the interfering noise, accurate channel estimates were obtained for both low-pass and multi-band-pass channels. When the highly correlated, nearly periodic noise samples were used, the channel estimates did not converge to the known channels. But the noise reduction in these cases was worsened only during times of changing noise characteristics. The filter readily adapted to new noise




Figure 4.3 AN EXAMPLE OF SIGNALS PROCESSED
 a) Original Signal b) Noisy Signal c) Signal Estimate




Figure 4.4 SPECTRA OF THE SIGNALS IN FIGURE 4.3
 a) Original Signal b) Noisy Signal c) Signal Estimate




Figure 4.5 TYPICAL ROOM CHANNELS

characteristics. Though the channel estimate did not converge to the known channels in these cases, the goal of noise reduction was still achieved.

Room Simulations

It was decided to predict the degree of cancellation possible in a hard-walled room about fifteen feet square. Estimates of its transfer function from a point near one wall to points in the room were made. The previously recorded noise was then digitally convolved with each of

the room estimates and two separate experiments were performed on the resulting signals.

In the first of these, the original noise signal was used as the reference noise, and one of the filtered signals was used as the noisy signal. No speech was added to this signal for this experiment. While the original room estimates were 4096 points in length, the adaptive filter was constrained to a length of 3000 points. An adaptation time-constant of approximately 0.4 seconds was specified. Noise reduction of -25 dB was measured for this experiment.

In the second experiment, the reference noise was provided by one of the digitally filtered signals, while the noisy signal was the output from the other room channel estimate. Again 3000 points were specified for the adaptive filter's length, half of them before $t = 0$. The resulting noise reduction measured was -12 dB.

Multichannel Experiments

The final group of synthetic experiments performed were the multichannel experiments. It is easy to conceive of a situation where the noise source is very broad-band and the noise received at the signal sensor exhibits similar characteristics, but the reference noise sensor, due to physical constraints, must be placed in a position where some segment of the noise spectrum has been greatly attenuated. In such a situation, it is imperative

that the channel be estimated accurately in all frequency bands, since the noise interfering with the signal has significant energy at all frequencies, even though the reference noise does not. Such a situation presents a particularly difficult challenge to the adaptive noise cancellation algorithm. Feeling that the problem could be minimized by the application of the multichannel processing scheme referred to in chapter 3, experiments were performed to evaluate the effectiveness of that technique in such a situation.

In order to perform such an evaluation, the filter shown in figure 4.6a was generated and used as the channel G_2 while G_1 was forced to be the identity system. The standard ANC algorithm was then applied, attempting to estimate G_2^{-1} . The reference noise was then divided into two frequency bands, and the multichannel processing technique applied. A total equivalent filter was then calculated for comparison with the estimate of the single channel scheme. A block analysis was also performed for comparison.

The standard ANC algorithm provided -1.5 dB noise reduction in this perverse environment, while the multichannel scheme attained a level of -35 dB noise reduction. Figures 4.6 and 4.7 show the channel G_2 and the estimates of its inverse obtained via the various procedures.




Figure 4.6 A CHANNEL AND THE ESTIMATES OF ITS INVERSE
 a) Test Channel (G_2) b) Block Processed Estimate
 c) Adaptive Estimate d) Multi-channel Estimate




Figure 4.7 SPECTRA OF CHANNELS IN FIGURE 4.6
 a) Test Channel (G_2) b) Block Processed Estimate
 c) Adaptive Estimate d) Multi-channel Estimate

Acoustically Recorded Experiments

Other experiments were performed in an actual acoustic environment. The digitized noise sources were played through a single multielement BOSE loudspeaker and digitally recorded through two separate SONY ECM-270 microphones placed at varying locations in the previously mentioned, hard-walled room. A speaker was situated near one of the microphones and spoke into it to provide the noisy signal, while the other microphone provided a signal which was used as the noise reference signal. For some problem evaluation work, the noise reference was picked up by electrically monitoring the speaker driving signal rather than through a second microphone. The active interval of the filters being estimated was arbitrarily chosen to be centered around $t = 0$. This was accomplished by delaying the noisy signal by half as many points as were specified as the length of the filter to be estimated.

A single channel room estimation was performed using the previously mentioned electrically monitored speaker driving signal as the reference noise, and a simultaneously acoustically recorded signal as the noisy signal. This corresponded to forcing G_2 in figure 2.2 to be an identity system. The noise reduction achieved in this experiment was -24 dB.

Many experiments were performed using one acoustically recorded signal as the reference noise and a

simultaneously recorded signal as the noisy signal. Varying filter lengths were specified and two methods of estimating the noise cancelling filter were employed. One set of experiments used the block analysis approach described in chapter 2 to derive a single stationary filter which was used in the noise canceller. The second set employed the standard ANC procedure.

The results of these experiments are compared in figure 4.8. The variation from record to record when using the block analysis method of filter determination was considerable. For some experiments, results degraded by up to 20 dB from those shown were observed in some records. Improvements of several dB in other records were not uncommon. The results shown, however were the improvement achieved in regions of little fluctuation, and correspond to the same records used for the adaptive filtering results. The results obtained with the adaptive filter were much more consistent from record to record except where they improved considerably (at least 10 dB for the longest filter length) during periods of highly correlated noise.




Figure 4.8 RESULTS OF ACOUSTICALLY RECORDED EXPERIMENTS
Filter Length vs. Noise Reduction in a Hard-walled Room
A--Block Analysis ©--Adaptive Noise Cancellation

CHAPTER 5

CONCLUSIONS

Observations on Results

The performance of the experiments presented in chapter 4 has provided much information and insight into the behavior of Adaptive Noise Cancellation. An analysis of the significance of the reported results is now presented.

The implementation of synthetic channels with short active intervals provided helpful understanding of the effects of deviation from the conditions assumed in mathematical analyses of the ANC process. When the reference noise measurements are not statistically independent, as they are not when the noise is nearly periodic, the adaptive filter is not guaranteed to converge to the optimal filter. Experimental examination of its performance under these conditions, however, showed that noise reduction was as good as, or better than that achieved when a white noise source was employed. Further examination of the adaptive filter's impulse response showed that such estimates of the channels were excellent in frequency bands where significant noise energy was present, and very poor where no noise was present. This

was not unexpected, since the adaptive filter's impulse response is a linear combination of previous reference noise sample vectors. The optimal filter was not unique for such noise. As the frequency of the noise was varied, however, degradation occurred because the filter was not an accurate channel estimate at all frequencies.

The error caused by adapting after achieving a perfect channel estimate emphasized the fact that there is a feedback path from the signal output to the channel estimate. This is evident from the description of the adaptive process found in equation 2.18. If rapid adaptation rates were chosen, this caused an effect much like an echo to be noticed at the signal estimate output.

A comparison of the results obtained from synthetic and actual rooms (-25 dB vs. -24 dB in the single channel case, and -12 dB vs. -10.5 dB for the two channel case) indicates that the assumption that a room can be modelled as a linear, stationary channel is valid. The near agreement of the results also implies that the introduction of spatially distributed noise sources, such as multi-element loudspeakers, is not a cause of great degradation in performance if such distribution is confined to a relatively small region.

Perhaps the most significant results reported are those produced by multichannel processing. The introduction of multichannel processing significantly improved

noise cancellation performance for the cases tested. The improved performance can be directly attributed to improved convergence of the adaptive filter, and specifically to the capability of the multichannel technique to allow different convergence rates in different bands. It is possible that the single channel technique would eventually converge to the optimal channel, or that a different adaptation rate would improve its performance. In either case, however, it would effectively converge in different frequency bands at different rates, and in the case of faster adaptation, a higher residual error would remain. When coupled with the increased filter length capacity available from multichannel processing (due to its inherent parallel processing nature) it is felt that its development has been significant.

The comparisons of filter length versus noise reduction for acoustically produced signals show relative performance losses caused by filter truncation. They are applicable to a single, hard-walled, acoustically live room about fifteen feet square, and ought not to be considered universally attainable levels. While the near agreement of the block analysis technique and the ANC algorithm indicates that the results are accurate, it is possible that extending the length of the data samples processed could result in changes to the reported results. That is, though the noise reduction levels reported

appeared to be relatively stable, the exceeding length of time required to perform the experiments on the long filter lengths due to the software implementation of the ANC algorithm on a general purpose computer (DEC PDP-10), precluded the use of exceedingly long data segments.

The absolute noise reduction obtainable for a given filter length is extremely dependent upon the physical environment where the process is being employed. For this reason, it is impractical to assume that the levels reported will be obtained in any particular situation. If the process is to be implemented in a less reverberant environment, however, considerably better performance would be expected.

The comparisons of the ANC approach and the global block analysis yield some surprises. The first of these is that the ANC approach consistently performed better than the supposedly optimal channel estimate. This was because the block estimate was calculated over the entire noise sample, which was nonstationary. It consisted of alternating segments of white gaussian noise and a highly correlated noise sample produced by a square-wave generator, as described previously. The block analysis not developed under such assumptions and attempted to minimize the total output energy. In doing so, it often allowed slightly poorer signal estimates for many blocks in return for one particularly good estimate for a shorter

period. In addition, the adaptive technique was not constrained to perform the estimation of a single channel, but was allowed to vary it to improve estimation on a point by point basis. And the second surprise is the anomaly seen for the 3000 point block estimate. This was the result of the failure of Levinson's recursive inversion algorithm [40] to obtain an accurate channel estimate.

Summary

Many experiments have been performed which indicate that Adaptive Noise Cancellation is capable of significantly reducing noise in audio signals if properly applied. The results of the research reported indicate that ANC techniques are promising for use in practical situations where audio noise levels are so high that other techniques are not applicable. The situation must be such that the additional costs of this type of processing can be justified. Requirements for two signal measurements must be justified and sufficient computing power must be available. For audio signal processing, some sort of automatic gain control must be provided. If parallel processing is to be used due to the length of the adaptive filter required, or if knowledge of the environment and noise source indicates that convergence problems are likely, implementation of the multichannel processing scheme presented appears to be a desirable, effective solution.

Future Work

Though the ANC technique has been shown to be effective for noise reduction in audio signals, there is still work to be done to understand it more fully and improve its performance. Some work must be done on channel estimation of proposed operating environments for ANC implementation. Such work could be coupled with further truncation error development to predict achievable performance for a given environment. Alternatively, using noise samples from such environments alone, or in conjunction with channel estimates, accurate simulations of such environments could be made to determine the required adaptive filter lengths necessary for attaining the desired noise reduction levels. Experimental work concerning development of band-pass filters for use with the multichannel process might also be performed to optimize performance and minimize observed ringing of the total equivalent filter's response in the cross-over bands.

APPENDIX A

CALCULATION OF THE CONDITIONAL MEAN

Let us assume that we have four jointly independent, zero mean Gaussian Random Variables, S , N , m_1 , and m_2 , and that their variances are known. Furthermore, let us assume that we can measure directly only the sums

$$(A.1) \quad x = S + N + m_1 \quad \text{and} \quad V = N + m_2$$

Then the mean of $f(S|X,V)$, where $f(\cdot)$ is the probability density of (\cdot) , is the MMSE estimate of S . Let us, then, calculate $E(S|X,V)$. Applying Bayes' Theorem

$$(A.2) \quad f(S|X,V) = \frac{f(X,V|S) f(S)}{f(X,V)}$$

The denominator of this expression serves only to normalize the distribution; the form, and hence the mean of the distribution, can be determined from the numerator alone. We may proceed directly to calculate the numerator. The conditional variance of X , given S , is simply $\sigma_N^2 + \sigma_{m_1}^2$. The covariance between X and V , given S , is

$$(A.3) \quad \text{COV}[X,V] = E\{(N + m_1)(N + m_2)\} = \sigma_N^2.$$

The covariance matrix of the augmented vector (X,V) is thus

$$(A.4) \quad R = \begin{bmatrix} \sigma_N^2 + \sigma_{m1}^2 & \sigma_N^2 \\ \sigma_N^2 & \sigma_N^2 + \sigma_{m2}^2 \end{bmatrix}$$

and the conditional mean vector is

$$(A.5) \quad M = (S, 0)^T .$$

The inverse of R is

$$(A.6) \quad R^{-1} = \frac{\begin{bmatrix} \sigma_N^2 + \sigma_{m2}^2 & -\sigma_N^2 \\ -\sigma_N^2 & \sigma_N^2 + \sigma_{m1}^2 \end{bmatrix}}{(\sigma_N^2 + \sigma_{m1}^2)(\sigma_N^2 + \sigma_{m2}^2) - \sigma_N^4}$$

$f(X, V, S)$ is proportional to

$$\exp\{-.5[X-S, V] R^{-1}[X-S, V]^T\} .$$

Since we are looking for

$$(A.7) \quad \hat{S} = E\{S|X, Y\} ,$$

let us find the terms involving S.

We find

$$\frac{((\sigma_N^2 + \sigma_{m2}^2)S - 2[(\sigma_N^2 + \sigma_{m2}^2)X - \sigma_N^2 V]S)}{2[\sigma_N^2(\sigma_{m1}^2 + \sigma_{m2}^2) + \sigma_{m1}^2\sigma_{m2}^2]}$$

to which we add (since we must multiply by $f(s)$)

$$- \frac{S}{2\sigma_S^2} .$$

If we then normalize the S term we are left with $-2\hat{S}$ as the coefficient to S, yielding

$$(A.8) \quad \hat{S} = \frac{\sigma_S^2 [(\sigma_N^2 + \sigma_{m2}^2)X - \sigma_N^2 V]}{\sigma_N^2(\sigma_{m1}^2 + \sigma_{m2}^2) + \sigma_{m1}^2 \sigma_{m2}^2 + \sigma_S^2(\sigma_N^2 + \sigma_{m2}^2)}$$

Thus we see that S is a linear combination of X and V.

$$(A.9) \quad S = aX + bV$$

where

$$(A.10) \quad a = \frac{(\sigma_N^2 + \sigma_{m2}^2)}{\sigma_N^2(\sigma_{m1}^2 + \sigma_{m2}^2) + \sigma_{m1}^2 \sigma_{m2}^2 + \sigma_S^2(\sigma_N^2 + \sigma_{m2}^2)}$$

and

$$(A.11) \quad b = -a \frac{\sigma_N^2}{\sigma_N^2 + \sigma_{m2}^2}$$

rewriting a

$$(A.12) \quad a = \frac{\sigma_S^2}{\sigma_S^2 + \left[\sigma_{m1}^2 + \sigma_{m2}^2 \left\{ \frac{\sigma_N^2}{\sigma_N^2 + \sigma_{m2}^2} \right\} \right]}$$

we see, by careful examination, that the least mean squares estimate of S is formed by first calculating a Wiener filtered estimate of N using only the known statistics and V. This estimate is then subtracted from X, and the result is filtered to minimize the effects of the remaining noise, including m1 and the residual from our misestimation of N. Had one of the measured signals contained not N, but a signal correlated with N, an additional coloration of V would have been required.

APPENDIX B

CHARACTERISTICS OF THE LMS ALGORITHM AND NOISE CANCELLATION

The LMS algorithm developed by Widrow and Hoff has been applied to the problem of noise cancellation in many contexts. Many of its properties have been written about frequently. Some of them are briefly described here. Most can be found described in more detail in reference [28].

In chapter 2, equations 2.19 and 2.20, it was shown that the channel estimate H converged to the optimal solution for $0 < \mu < \frac{1}{\lambda_{\max}}$. This condition is related to the total reference noise input power by

$$\lambda_{\max} < \text{TR}[R] = \sum_{i=1}^L E\{v_{n+1}^2\}.$$

If we assume ergodicity and compute the expected value as a time-average, we see that the input power is an upper bound on the maximum eigenvalue of R . It is the sum of all of the eigenvalues. Since R is positive definite, each of the eigenvalues must be positive. In the special case where all of the eigenvalues are equal, such as when R represents truly white, stationary noise, this bound is larger than necessary by a factor of L , where L is the dimension of the vector V .

The weights of the adaptive filter adapt in such a way that their time-history is the sum of L complex exponentials with time-constants

$$\tau_i = \frac{1}{2\mu\lambda_i} \quad i = 1, 2, \dots, L$$

where the λ_i are the eigenvalues of R . If they are all equal,

$$\tau = \frac{1}{2\mu\lambda}$$

Since the mean squared error is a quadratic function of the weight values, the time-constant representing the rate at which the mean squared error approaches its final value is

$$\tau_{mse} = \frac{\tau}{2}$$

If perfect noise cancellation cannot be achieved, a "misadjustment" measure defined as the ratio of the excess mean squared error to the minimum mean squared error has been defined. It has been shown to be

$$M = \frac{1}{2} \sum_{i=1}^L \frac{1}{\tau_i}$$

for statistically stationary processes. It is proportional to the number of weights, L , and inversely proportional to the time constant. From equation 2.6 we see that the optimal noise estimate is

$$u_n = \sum_{i=-\infty}^{\infty} h_i * v_{n-i} .$$

In practice we could only hope for

$$u_n = \sum_{i=N}^{N+L-1} h_i * v_{n-i}$$

and we can actually get

$$u_n = \sum_{i=N}^{N+L-1} \hat{h}_i * v_{n-i}$$

showing that we have two major sources of error. The first source is the error in the estimation of the h . This error can be decreased by increasing the time constant of adaptation (decreasing the adaptation constant μ) or by decreasing the number of taps in the active interval of the filter. It is this error which gives rise to the misadjustment mentioned.

The second source of error is the truncation of the infinite summation. We, thereby, assume that all truncated h are zero. If they are not, increasing the active interval of the filter (adding more terms) may result in a better noise estimate. It should be noted that increasing the number of terms indefinitely will not result in increasingly accurate noise estimates. As the truncation error decreases, the error in parameter estimation increases, and the misadjustment error increases. Of course, a correspondingly longer time-constant can be specified in order to reduce the misadjustment error.

REFERENCES

- [1] R. McAulay, "Optimum Classification of Voiced Speech, Unvoiced Speech and Silence in the Presence of Noise and Interference," Lincoln Lab Technical Note, 1976-7, June 1976.
- [2] N. Wiener, "Extrapolation, Interpolation and Smoothing of Stationary Time Series, with Engineering Applications," New York: Technology Press and Wiley, 1949.
- [3] R. Kalman, "On the General Theory of Control," in Proc. 1st IFAC Congress. London: Butterworth, 1960.
- [4] J. S. Lim, "Evaluation of Autocorrelation Subtraction Method for Enhancing Speech Degraded by Additive White Noise," submitted to IEEE Trans. on Acoustics, Speech, and Signal Processing.
- [5] N. J. Miller, "Removal of Noise From a Voice Signal by Synthesis," Ph.D. dissertation, University of Utah, 1973, rep. UTEC-CSc-74-013, May 1973.
- [6] B. Yegnanarayana, "Effect of Noise and Distortion in Speech on Parametric Extraction," 1976 IEEE International Conf. on Acoustics, Speech, and Signal Processing Record, pp. 336-339, April 1976.
- [7] R. Riegler and R. Compton, Jr., "An Adaptive Array for Interference Rejection," Proceedings of the IEEE, vol. 61, pp. 748-758, June 1973.
- [8] B. Widrow, et al., "Adaptive Antenna Systems," Proc. IEEE, vol. 55, pp. 2143-2159, Dec. 1967.
- [9] R. W. Lucky, "Automatic Equalization for Digital Communication," The Bell System Technical Journal, vol. 44, pp. 547-588, April 1965.
- [10] M. M. Sondhi, "An Adaptive Echo Canceller," The Bell System Technical Journal, vol. 46, pp. 497-511, March 1967.

- [11] K. H. Mueller, "A New Digital Echo Canceled for Two-wire Full-Duplex Data Transmission," IEEE Transactions on Communications, vol. COM-24, pp. 956-962, September 1976.
- [12] Stephen B. Weinstein, "Echo Cancellation in the Telephone Network," IEEE Communications Society Magazine, vol. 15, pp. 91-5, January 1977.
- [13] J. R. Rosenberger and E. J. Thomas, "Performance of an Adaptive Echo Canceller Operating in a Noisy, Linear Time-Invariant Environment," Bell System Technical Journal, vol. 50, pp. 785-813, March 1971.
- [14] J. Kaunitz, "Adaptive Filtering of Broadband Signals as Applied to Noise Cancelling," Ph.D. dissertation, Stanford University, 1972.
- [15] J. R. Glover, "Adaptive Noise Cancelling of Sinusoidal Interference," Ph.D. dissertation, Stanford University, 1975.
- [16] B. S. Atal and S. L. Hanauer, "Speech Analysis and Synthesis by Linear Prediction of the Speech Wave," J. Acoustical Soc. of America, vol. 50, pp. 637-655, Aug. 1971.
- [17] J. Makhoul, "Spectral Linear Prediction: Properties and Applications," IEEE Transactions on Acoustics, Speech, and Signal Processing, vol. ASSP-23, pp. 283-286, June 1975.
- [18] Thomas Kailath, "A View of Three Decades of Linear Filtering Theory," IEEE Transactions on Information Theory, vol. IT-20, pp. 146-181, March 1974.
- [19] S. F. Boll, "Suppression of Noise in Speech Using the SABER Method," technical report UTEC-CSc-78-073, Computer Science Department, University of Utah, April 1978.
- [20] S. F. Boll, "Improving Linear Prediction Analysis of Noisy Speech by Predictive Noise Cancellation," 1977 IEEE International Conf. of Acoustics, Speech, and Signal Processing Record, pp. 10-12, May 1977.
- [21] M. A. Atashroo, "Pole-Zero Modeling and its Application to Speech Processing," Ph.D. dissertation, University of Utah, 1976.

- [22] R. W. Christiansen, "Word Recognition in Continuous Speech Using Linear Prediction Analysis," Ph.D. dissertation, University of Utah, 1976.
- [23] H. W. Bode and C. E. Shannon, "A Simplified Derivation of Linear Least Square Smoothing and Prediction Theory," Proceedings of the IRE, vol. 38, pp. 417-425, April 1950.
- [24] G. E. Box and G. M. Jenkins, Time Series Analysis, Forecasting and Control, Holden-Day, San Francisco, 1976.
- [25] H. Robbins and S. Monroe, "A Stochastic Approximation Method," Annals of Math Statistics, vol. 22, pp. 400-407, 1951.
- [26] D. J. Sakrison, "Stochastic Approximation: A Recursive Method for Solving Regression Problems," in Advances in Communication Theory, vol. 2, Academic Press, New York, 1966.
- [27] B. Widrow and M. Hoff, Jr., "Adaptive Switching Circuits," in IRE WESCON Conv. Rec., Pt. 4, pp. 96-104, 1960.
- [28] Bernard Widrow, "Adaptive Filters," from Aspects of Network and System Theory, edited by R. E. Kalman and N. DeClaris, Holt, Rinehart and Winston, Inc., N.Y., 1970.
- [29] K. D. Senne, "Adaptive Linear Discrete-Time Estimation," Ph.D. dissertation, Stanford University, 1968.
- [30] Bernard Widrow, et al., "Adaptive Noise Cancelling: Principles and Applications," Proceedings of the IEEE, vol. 63, pp. 1692-1719, December 1975.
- [31] T. P. Daniell, "Adaptive Estimation with Mutually Correlated Training Samples," Ph.D. dissertation, Stanford University, 1968.
- [32] J. K. Kim and L. D. Davisson, "Adaptive Linear Estimation for Stationary M-dependent Processes," IEEE Transactions on Information Theory, vol. IT-21, pp. 23-31, January 1975.
- [33] B. Widrow, et al., "Stationary and Nonstationary Learning Characteristics of the LMS Adaptive Filter," Proceedings of the IEEE, vol. 64, pp. 1151-1162, August 1976.

- [34] J. L. Moschner, "Adaptive Filter With Clipped Input Data," SEL-70-053, Stanford Electronics Laboratories, June 1970.
- [35] T. M. McSherry, "A General Steepest Descent Algorithm," IEEE Transactions on Aerospace and Electronics Systems, vol. AES-12, pp. 12-22, January 1976.
- [36] Jin-Ichi Nagumo and Atsuhiko Noda, "A Learning Method for System Identification," IEEE Transactions on Automatic Control, vol. AC-12, no. 3, pp. 282-287, June 1967.
- [37] R. D. Gitlin, J. Mazo, and M. G. Taylor, "On the Design of Gradient Algorithms for Digitally Implemented Filters," IEEE Transactions on Circuit Theory, vol. CT-20, pp. 125-136. March 1973.
- [38] O. L. Frost, III, "Adaptive Least Squares Optimization Subject to Linear Equality Constraints," SEL-70-055, Stanford Electronics Laboratories, Aug. 1971.
- [39] R. Penrose, "On the Best Approximate solutions of Linear Matrix Equations," Proc. Cambridge Phil. Soc., vol. 52, pp. 17-19, 1956.
- [40] N. Levinson, "The Wiener RMS Error Criterion in Filter Design and Prediction," Journal of Mathematics and Physics, vol. 25, No. 4, pp. 261-278, 1947.
- [41] S. F. Boll, E. Ferretti, T. Petersen, "Improving Synthetic Speech Quality Using Binaural Reverberation," Conference Report of the 1976 IEEE ICASSP, pp. 705-508, 1976.